# A Nearly Optimal Contextual Bandit Algorithm

Mohammadreza Mohaghegh Neyshabouri, Kaan Gokcesu, Selami Ciftci
and Suleyman S. Kozat, *Senior Member, IEEE*

arXiv:1612.01367v1 [cs.LG] 5 Dec 2016

***Abstract*—We investigate the contextual multi-armed bandit problem in an adversarial setting and introduce an online algorithm that asymptotically achieves the performance of the best contextual bandit arm selection strategy under certain conditions. We show that our algorithm is highly efficient and provides significantly improved performance with a guaranteed performance upper bound in a strong mathematical sense. We have no statistical assumptions on the context vectors and the loss of the bandit arms, hence our results are guaranteed to hold even in adversarial environments. We use a tree notion in order to partition the space of context vectors in a nested structure. Using this tree, we construct a large class of context dependent bandit arm selection strategies and adaptively combine them to achieve the performance of the best strategy. We use the hierarchical nature of introduced tree to implement this combination with a significantly low computational complexity, thus our algorithm can be efficiently used in applications involving big data. Through extensive set of experiments involving synthetic and real data, we demonstrate significant performance gains achieved by the proposed algorithm with respect to the state-of-the-art adversarial bandit algorithms.**

***Index Terms*—Contextual bandits, universal, online learning, individual sequence manner, big data.**

**EDICS Category: MLR-SLER, MLR-APPL, MLR-LEAR.**

## I. Introduction

### A. Preliminaries

In this paper, we study contextual multi-armed bandit problem, which is a widely studied topic in the sequential learning theory [1]–[3]. In the classic multi-armed bandit setting there are $M$ actions, each of which is considered as a bandit arm. At each round, we choose one arm and obtain a reward (or loss) according to the selected arm, where the reward (or loss) of all the other arms are oblivious to us. The objective is to maximize the cumulative reward of the selected arms in a series of rounds. Since the reward we would obtain from the other arms remain hidden, this setting can be considered as a limited feedback version of the well studied problem of prediction with expert advice [4]–[7]. The multi-armed bandit problem naturally addresses the fundamental trade-off between exploration and exploitation in sequential learning theory [8]. Here, we must balance exploitation of actions that gave the highest payoffs in the past and exploration of actions that might give higher payoffs in the future. The multi-armed bandit problem has attracted significant attention due to applicability of the bandit setting in a wide range of real world learning applications from online advertisement [9] and recommender systems [10], [11] to clinical trials [12] and cognitive radio [13], [14]. As an example, in the online advertisement application, different ads available to display to users are modeled as the bandit arms and the act of clicking by the user on the displayed ad is modeled as the reward [9].

In a wide variety of applications of the bandit algorithms, we can obtain side information on the problem at hand before selecting the arms at each round [15], e.g., in the clinical trials the side information can be the age or the gender of the patient [16]. However, most of the conventional bandit algorithms do not exploit or fail to optimally exploit this information to improve their performance [17]–[19]. To remedy this, contextual multi-armed bandit algorithms are introduced as practical alternatives to simple multi-armed bandit methods [8], [9], [20]. In the contextual multi-armed bandit setting, the side information is represented as a context vector. As an example, in the online advertisement applications this context vector may contain certain information about the users such as historical activities, demographic information, etc. Here, the algorithm learns how to effectively use the side information vector to optimize its arm selection strategy and gain more rewards.

We study the contextual multi-armed bandit problem in an online setting, where we operate sequentially on a stream of observations from a possibly nonstationary, chaotic or even adversarial environment. In our framework, we have no statistical assumptions on the context vectors and behavior of the bandit arms so that our result are guaranteed to hold in an individual sequence manner [8]. We emphasize that we refrain from making any statistical assumptions since in most real life applications of the bandit algorithms, the underlying application can be highly nonstationary, e.g., in online ad selection [21], or may be even adversarial, e.g., in gambling [22].

To this end, we investigate the contextual multi-armed bandit problem from a competitive algorithm perspective [8]. Since we have no statistical assumptions on the context vectors and the bandit arms behavior, we define our performance with respect to a competition class of context dependent bandit arm selection policies. As the competition class, we use the class of all predetermined mappings from the space of context vectors to the bandit arms. Here, each mapping arbitrarily partitions the space of context vectors into several disjoint regions and maps each one of these regions to one of the bandit arms, i.e., at each round it selects the bandit arm corresponding to the region containing the observed context vector. A sample mapping from the space of context vectors to the bandit arms in a 2-armed bandit setting is shown in Fig. 1a. In this example,

This work is supported in part by Turkish Academy of Sciences Outstanding Researcher Programme, TUBITAK Contract No. 113E517.

M. Mohaghegh N., K, Gokcesu and S. S. Kozat are with the EE Department of Bilkent University, Ankara, Turkey, {mohaghegh, gokcesu, kozat}@ee.bilkent.edu.tr, tel: 0312 2902336. Selami Ciftci is with Turk Telekom Communications Services Inc., Istanbul, Turkey, selamiciftci@turktelekom.com.tr

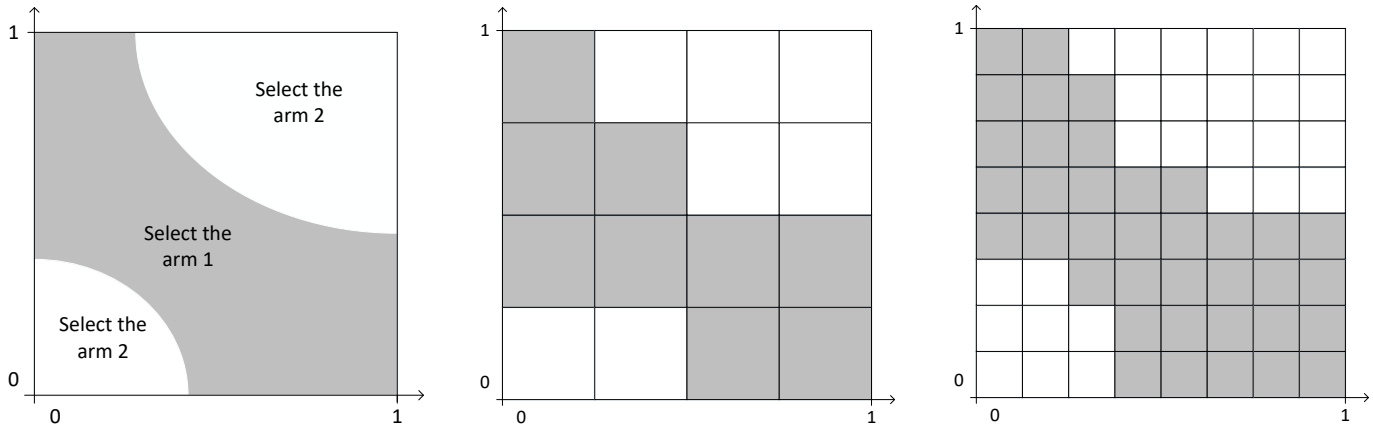


(a) An example mapping from the context space $[0,1]^2$ to the bandit arms set $\{1,2\}$.

(b) Closest mapping in the quantized competition class with 16 quantization levels to the mapping in Fig. 1a.

(c) Closest mapping in the quantized competition class with 64 quantization levels to the mapping in Fig. 1a.

Fig. 1: An example mapping from the context space to the arms set and its approximations in the quantized competition classes. In each mapping shown in the figure, the dark and bright sections are mapped to the arms 1 and 2, respectively.

the space of the context vectors is the 2-dimensional space $[0,1]^2$. We emphasize that the optimal mapping is the one with the optimal partition of the space of context vectors and also the best arm selection over all regions of this partition. However, such a mapping can only be known after we observe all the bandit arm losses both in the past and in the future, which is clearly impossible in the online setting. In this paper, our goal is to achieve the performance of this optimal mapping, which can optimally partition the space of the context vectors and assign the best arm to each region in hindsight before we even start processing or observe the data.

For this purpose, we quantize the space of the context vectors into a large number of disjoint regions and seek to achieve the performance of the best mapping from these quantized regions to the bandit arms. We call the class of all mappings from the quantized regions to the bandit arms as the quantized competition class. Note that as the number of quantization levels gets large enough, the optimal mapping in the quantized competition class gets as close as desired to the truly optimal mapping, hence we achieve the performance of the optimal mapping from the space of context vectors to the bandit arms. As an example, consider Fig. 1 again. This time suppose that the mapping in Fig. 1a is truly the optimal mapping from the space of context vectors to the bandit arms. In this case, the optimal mapping in the quantized competition classes for the number of quantization levels equal to 16 and 64 are shown in Fig. 1b and Fig. 1c, respectively. These figures demonstrate that if the number of quantization levels is large enough, the best mapping in the quantized competition class is close enough to the truly optimal mapping.

There are other algorithms such as *S-EXP3* [8] that quantize the space of context vectors into several disjoint regions and run independent non-contextual bandit algorithms over each one of the regions. If there is an infinitely large number of data, then *S-EXP3* achieves the performance of the optimal mapping in the quantized competition class. However, since we are working in an online setting, *S-EXP3* is prone to over fitting, which severely degrades its performance as the number of quantization levels increases. To mitigate such difficulties, we use a notion of binary tree to quantize the space of context vectors into disjoint regions in a nested structure. Using the hierarchical properties of this tree, we combine all possible mappings from the quantized regions to the arms set in a mixture of experts setting [23]. We carefully give specific initial weights to the mappings such that the mappings that use finer partitions of the space of context vectors get higher initial weights in comparison to the mappings using coarser partitions. Using this initial weighting, at the beginning rounds, when we do not have access to much information, we stick to the mappings over coarser partitions and as the time goes on and we gain more information, the effect of the initial weights disappear such that we finally achieve the performance of the optimum mapping from the finest partition to the bandit arms.

The single tree-based combination of mappings described in the previous paragraph needs to know the number of disjoint regions in the partition corresponding to the optimal mapping in order to tune its parameter. In order to remove this dependency, we run a specific number of copies of our basic algorithm with carefully selected parameters in parallel and adaptively combine them using the well known *EXP4* [8] algorithm. We show that using this combination, while retaining the order of space and computational complexities, we achieve the performance of the optimal mapping from the quantized regions to the arms set without any a priori information about this mapping.

Finally, we introduce an efficient quantization method for the generic case of any $n$-dimensional context space and show that if the arm losses are Lipschitz continuous functions of the context vectors at each specific round, using the proposed quantization method, our algorithm asymptotically achieves the performance of the optimal mapping from the context space to the bandit arms (with the optimal partition of the context space and the best arm assignments to the different regions of this partition) as the depth of our tree (or equivalently

the number of quantization levels) increases. We emphasize that the assumption on the loss functions being Lipschitz continuous with respect to the context vectors does not have any conflict with the adversarial setting of the problem. The loss functions can be quite different from a round to the next round as long as they are Lipschitz with respect to the context vectors at each round.

### B. Prior Art

The contextual bandit problem is mostly studied in the stochastic setting [20], [24]–[26], where it is assumed that at each round, first a pair of context vector and loss vector (including losses corresponding to all arms) is drawn from an unknown distribution, then, the context vector is shown to the user and an arm is selected and finally the loss of the selected arm is revealed. Further studies are done with additive assumptions on specific relationships between the context vectors and the arm losses, from a linear relation as in [9] and [27] to a more generic class of relations in [28].

An alternative to the stochastic setting for contextual bandit problem is the adversarial setting, where we refrain from any assumptions on the behaviour of the context vectors and bandit arms. Well known *S-EXP3* and *EXP4* algorithms [8] approach the contextual bandit problem in the adversarial setting. These algorithms can be interpreted as two mixture-of-expert type algorithms which combine different mappings from a quantized context space to the bandit arms and are proved to have performance guarantees of $O(\sqrt{TN})$, where $T$ is the number of rounds and $N$ is the number of quantization levels. We, for the first time in the literature, introduce a contextual bandit algorithm, which combines a large number of mappings from the quantized space of context vectors to the arms set and achieves a performance guarantee of $O(\sqrt{T \ln N})$. Our algorithm uses a notion of context tree to keep the space and computational complexities affordable.

The context trees are widely used in different applications including but not limited to data compression [29], [30], estimation [31], [32], communications [33], regression [34], [35] and classification [36]. In all aforementioned applications, context trees are used to partition the context space in a nested structure, run an independent adaptive model over each one of the tree nodes and combine the models to achieve the performance of the best context dependent model in a computationally efficient manner. In all of the conventional applications, a binary tree of depth $D$ is used to combine approximately $1.5^{2^D}$ context dependent adaptive models and achieve the performance of the best model in the combination. We use a novel notion of a binary tree designed for the contextual bandit setting, which directly combines $2^{2^D-1}$ fixed (not adaptive) context dependent arm selection models, where each model consists of a specific partition of the space of context vectors and a bandit arm assigned to each region of this partition. The main novelty in our binary tree is that we directly combine a larger class of fixed models instead of running a relatively small number of adaptive models as in the conventional binary trees. Using such a binary tree we can easily optimize the parameter of our algorithm as opposed to the normal binary tree based algorithms.

### C. Contributions

- We introduce novel and efficient contextual bandit arm selection algorithms, which achieve a performance as close as desired to the performance of the best mapping from the space of the context vectors to the bandit arms as the number of quantization levels increases.
- We introduce an efficient quantization method and show that using this quantization method, our algorithm asymptotically achieves the performance of the optimal fixed mapping from the space of the context vectors to the arms set (in the average loss per round sense) as the number of quantization levels increases.
- We propose an efficient implementation of our algorithms such that we achieve this performance with a computational complexity only log-linear in the number of quantization levels.
- We demonstrate the significant performance gain of the proposed algorithms in comparison to the state-of-the-art contextual bandit algorithms through several experiments involving both synthetic and real data.

### D. Organization of the Paper

The paper is organized as follows. In Section II, we describe the contextual multi-armed bandit framework. We explain a brute force mixture of experts approach and its challenges in Section III. We explain the notion of binary trees in Section IV. Then, in Section V, we implement our algorithm using this binary tree. This algorithm needs to know a specific parameter of the optimal mapping in our mixture to achieve its performance. In Section VI, we remove this requirement and introduce our final algorithm, which achieves the performance of the optimal mapping in the mixture without any a priori information. We introduce an efficient quantization method in Section VII, and show that using this quantization method our algorithm is competitive against any mapping from the context space to the bandit arms. Section VIII contains the simulation results over several synthetic and well known real life datasets followed by the concluding remarks in Section IX.

## II. Problem Description

In this paper, all vectors are column vectors and denoted by boldface lower case letters. For a $K$-element vector $\mathbf{u}$, $u_i$ represents the $i^{\text{th}}$ element and $\|\mathbf{u}\| = \sqrt{\mathbf{u}^T\mathbf{u}}$ is the $l^2$-norm, where $\mathbf{u}^T$ is the ordinary transpose. We show the indicator function by $\mathbf{1}_{\{\text{condition}\}}$, which is equal to 1 if the condition holds and 0 otherwise. We say that a function $f(.) : \mathbb{R}^n \to \mathbb{R}$ is Lipschitz continuous with Lipschitzness constant $c$ over the region $W \subset \mathbb{R}^n$, if there is a (necessarily nonnegative) constant $c$ such that for all $\mathbf{x}_1, \mathbf{x}_2 \in W$, $|f(\mathbf{x}_1) - f(\mathbf{x}_2)| \leq c\|\mathbf{x}_1 - \mathbf{x}_2\|$.

We study the contextual multi-armed bandit problem in an adversarial setting. In the multi-armed bandit problem, we have $M$ bandit arms $\{1, ..., M\}$ and at each round $t$, we select one of the bandit arms $I_t \in \{1, ..., M\}$, and based on this selection, we suffer a loss $l_{t,I_t}$. We do not observe the loss of the unchosen arms. The objective is to minimize the accumulated loss incurred in a series of $T$ rounds, i.e,

$\sum_{t=1}^{T} l_{t,I_t}$ (where $T$ is unknown in our framework). In the contextual version of the multi-armed bandit problem at each round $t$, before selecting our bandit arm we observe a variable $\mathbf{s}_t$, called the context vector, from a set of possible context vectors $S$, called the context space. As an example in Fig. 1, $S$ is $[0,1]^2$ and each $\mathbf{s}_t$ corresponds to a vector in this space. For notational simplicity, we assume $l_{t,I_t} \in [0,1]$, however, it can be straightforwardly shown that our results hold for any bounded loss after shifting and scaling in magnitude. We work in the adversarial setting such that we do not assume any statistical model on the behavior of the context vectors and the bandit arms [22], hence our algorithms are guaranteed to work in an individual sequence manner. Our algorithms are strictly sequential such that at each round they selects an arm $I_t$ according to the information coming from previous rounds including observed context vectors, selected arms and their corresponding loss, alongside the context vector we are currently observing, i.e.,

$$I_t = f(\mathbf{s}_t; \mathbf{s}_{t-1}, I_{t-1}, l_{t-1,I_{t-1}}; ...; \mathbf{s}_1, I_1, l_{1,I_1}). \quad (1)$$

Consider a deterministic mapping $g(\cdot)$ from the context space to the arms set, i.e., $g(\cdot) : S \to \{1, ..., M\}$. Each such mapping is composed of a partition of the context space and an arm assigned to each region of this partition. An example mapping is shown in Fig. 1a for the case when the context space is $S = [0,1]^2$ and there are 2 bandit arms. In this figure, the dark and bright regions are mapped to the first and second arms, respectively. Each such mapping $g(\cdot)$ can be seen as a bandit arm selection policy, which at each round $t$, gets the context vector $\mathbf{s}_t$ as an input and outputs the bandit arm determined by $g(\mathbf{s}_t)$. In the following, we quantify the mappings by the number of disjoint regions they have in their partition of the context space. In our framework, the selected arms within neighbor disjoint regions cannot be the same since in this case neighbor regions can be merged to form one region.

Since we have no statistical assumptions on the behavior of the context vectors and the loss of the bandit arms [22], we define our performance with respect to the optimum (minimum loss) mapping in a specific class of mappings called the competition class. We use the notion of regret to define our performance against the competition class $U$ as

$$\mathcal{R}(T,U) \triangleq \max_{g \in U} \mathbb{E}\left[\sum_{t=1}^{T} l_{t,I_t} - \sum_{t=1}^{T} l_{t,g(\mathbf{s}_t)}\right], \quad (2)$$

where we denote the regret against the mappings in $U$ accumulated in $T$ rounds as $\mathcal{R}(T,U)$. We define $C$ as the class of all arbitraty mappings from the context space to the arms set, i.e., $C = \{g(\cdot) : S \to \{1, ..., M\}\}$. We emphasize that the ultimate goal is to achieve the performance of the best mapping in $C$ (in the average loss per round sense).

In the next section, we explain a brute force mixture of experts approach that constructs a large class of mappings from the space of context vectors to the bandit arms and adaptively combines them to achieve the performance of the best mapping in the class.

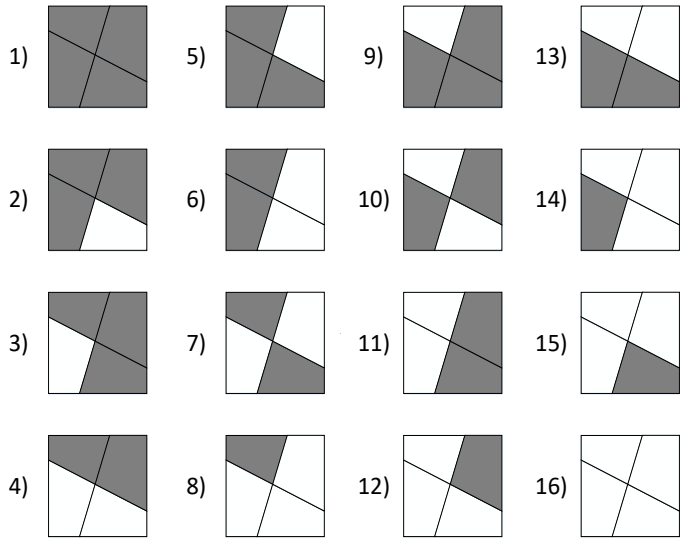


Fig. 2: All possible mappings in a 2-armed bandit problem with a predetermined quantization of the context space $S = [0,1]^2$ into 4 regions. In each mapping shown in the figure, the dark and bright regions are mapped to the arms 1 and 2, respectively.

## III. A Contextual Bandits Algorithm Based on Mixture of Experts

The ultimate goal in the contextual bandit problem is to achieve the performance of the optimal mapping from the context space to the bandit arms (with optimal partition of $S$ and optimal assignment of bandit arms to each region). Since the class of all arbitrary mappings, i.e., $C$, is too powerful to compete against without knowing the optimal partition of the context space, as the first step we quantize the context space $S$ into $N$ disjoint regions $r_1, r_2, ..., r_N$. We denote the set of these $N$ regions by $S^N = \{r_1, r_2, ..., r_N\}$ and call each mapping from $S^N$ to $\{1, ..., M\}$, i.e., $g(\cdot) : S^N \to \{1, ..., M\}$, as a quantized mapping of level $N$. Two examples of such quantized mappings of different levels for the case of 2-armed bandit with the context space of $[0,1]^2$ are shown in Fig. 1b and Fig. 1c. For each different mapping in this figure, the dark and bright regions are mapped to the first and second arms, respectively. We define $C^N$ as the class of all quantized mappings of level $N$, i.e., $C^N = \{g(\cdot) : S^N \to \{1, ..., M\}\}$, and seek to achieve the performance of the best quantized mapping in $C^N$. Intuitively, the performance of the best quantized mapping in $C^N$ gets closer to the performance of truly optimal mapping in $C$ as $N$ increases. As an example, suppose that the mapping shown in Fig. 1a is the optimal mapping in $C$. In this case the mappings in Fig. 1b and Fig. 1c will be the best mappings in $C^{16}$ and $C^{64}$, respectively.

Consider a contextual bandit problem with $M$ arms, where we have quantized the context space into $N = 2^D$ regions, where $D$ is an integer. Since we have $M$ arms, we define $M^N$ different experts $E_j, 1 \leq j \leq M^N$, corresponding to different assignment of arms to different regions, where each expert follows one of the mappings in $C^N$. An example of all 16 mappings followed by the experts for the case of $M = 2$

and $N = 4$, is shown in Fig. 2, where unlike Fig. 1 we choose a nonuniform quantization without loss of generality. In this figure the regions which are mapped to the first and second arms are depicted in gray and white, respectively. Each of these experts, i.e., a specific partition of $[0,1]^2$ and specific assignment of bandit arms to this partition, suggest a specific bandit arm to select (according to the observed context vector) at each round $t$. One of these experts is optimal for the underlying sequence of experts, however, naturally we do not know which. Hence, instead of committing to a single expert, we next use a mixture of experts notion.

In order to achieve the performance of the best expert, we assign each of the experts a weight (which shows our trust on the suggestion of this expert) and use a weighted version of *EXP4* algorithm [8] to adaptively combine the experts. The weight of the $i^{\text{th}}$ expert at round $t$ is denoted by $\alpha_{t,i}$. At each round $t$, after observing $\mathbf{s}_t$, we randomly select one of the experts using the probability simplex $\boldsymbol{\beta}_t = (\beta_{t,1}, ..., \beta_{t,M^N})$, where $\beta_{t,j} = \alpha_{t,j} / \sum_{k=1}^{M^n} \alpha_{t,k}$ is the normalized weight of the $j^{\text{th}}$ expert. Note that the probability of selecting each arm follows the probability simplex $\mathbf{p}_t = (p_{t,1}, ..., p_{t,M})$, where

$$p_{t,i} = \sum_{j=1}^{M^N} \beta_{t,j} \mathbf{1}_{\{g_j(\mathbf{s}_t)=i\}}. \tag{3}$$

We set initial values of the expert weights, i.e., $\alpha_{1,i}$'s, according to our a priori trust on the experts and use the *EXP4* algorithm [8] to update these weights according to their performance such that at each round $t \geq 2$, we have

$$\alpha_{t,i} = \alpha_{1,i} e^{-\eta \sum_{\tau=1}^{t-1} \tilde{l}_{\tau,g_i(\mathbf{s}_\tau)}}, \tag{4}$$

where $\eta \in \mathbb{R}^+$ is a constant called the learning rate and $\tilde{l}_{\tau,g_i(\mathbf{s}_\tau)}$ is the unbiased estimator of $l_{\tau,g_i(\mathbf{s}_\tau)}$, where $g_i(.)$ is the mapping followed by the $i^{\text{th}}$ expert $E_i$. Since we do not observe the loss of the unchosen arms, we have to use this unbiased estimator

$$\tilde{l}_{t,m} = \begin{cases} \frac{l_{t,m}}{p_{t,m}} & m = I_t \\ 0 & m \neq I_t \end{cases}, \tag{5}$$

which gives $\mathbb{E}[\tilde{l}_{t,m}] = l_{t,m}$, i.e., the expected value of each estimate is equal to its true value. Using this bandit arm selection probability assignment given in (3), (4), (5), we have the following regret result.

**Theorem 1.** *Consider an $M$-armed contextual bandit problem. Given that the context space is quantized into $N$ disjoint regions, Let $E_j$ be an expert which follows one of the $M^N$ possible mappings from the quantized context space to the bandit arms. Following the method detailed in Section III, $\mathcal{R}(T, E_j)$ satisfies*

$$\mathcal{R}(T, E_j) \leq \frac{\ln(1/\beta_{1,j})}{\eta} + \frac{MT\eta}{2}, \tag{6}$$

*where $T$ is the number of rounds, $\eta \in \mathbb{R}^+$ is the learning rate parameter used in* (4) *and $\beta_{1,j}$ is the normalized initial weight of the $j^{th}$ expert $E_j$.*

The proof of Theorem 1 follows the same steps as the proof of Theorem 4.2 in [8] with small variations due to our arbitrary initial weighting as opposed to uniform initial weights of the experts in [8]. The proof is provided in Appendix for completeness.

We point out that in the basic form of *EXP4* algorithm [8], the initial weights of the experts are equal, i.e., $\beta_{1,j} = 1/(M^N)$ for $j = 1$ to $M^N$, which yields a regret upper bound of $\sqrt{(NMT \ln M)/2}$, with optimum selection of $\eta = \sqrt{(2N \ln M)/MT}$. Furthermore, *S-EXP3* algorithm [8] achieves a regret upper bound of the same order ($\sqrt{2NMT \ln M}$) using an independent *EXP3* algorithm over each quantized region of the context space. Since we want to increase the number of quantization levels, this square root dependency of the regret bound on the quantization level is highly problematic. On the other hand, working with these $M^N$ parameters $\alpha_{t,1}, ..., \alpha_{t,M^N}$ has quite high space and computational complexities of $O(M^N)$. In the next section, we introduce a binary tree notion. Later, we use this tree to implement the described combination with an intelligent initial weighting (which leads to a regret upper bound with logarithmic dependency on the number of quantization levels $N$) and significantly low space and computational complexities.

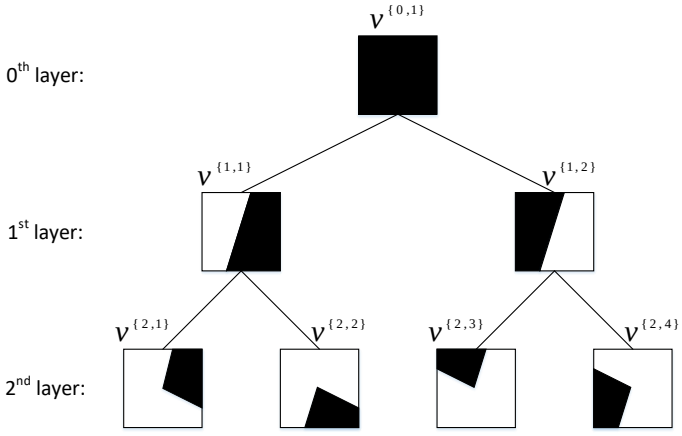


Fig. 3: A binary tree of depth $D = 2$, over the context space $[0,1]^2$. The regions corresponding to each node are filled with black color.

## IV. Binary Tree

We use a binary tree of depth $D$ to quantize the context space into $N = 2^D$ disjoint regions in a nested structure. As an example, consider the binary tree of depth 2 in Fig. 3, which quantizes the 2-dimensional context space $S = [0,1]^2$. Each node of such binary tree corresponds to a region of the context space, as shown in the figure. The region corresponding to each node, is the union of regions of its child nodes. We define $v^{\{i,j\}}$ as the $j^{\text{th}}$ node at $i^{\text{th}}$ layer of the tree, where $0^{\text{th}}$ layer is the top layer. We also define $R^{\{i,j\}}$ as the region corresponding to $v^{\{i,j\}}$. Note that using this notation we have $j \leq 2^i$ for all nodes. In the sequel, we use this binary tree to compactly represent our experts and combine them in a computationally efficient manner.

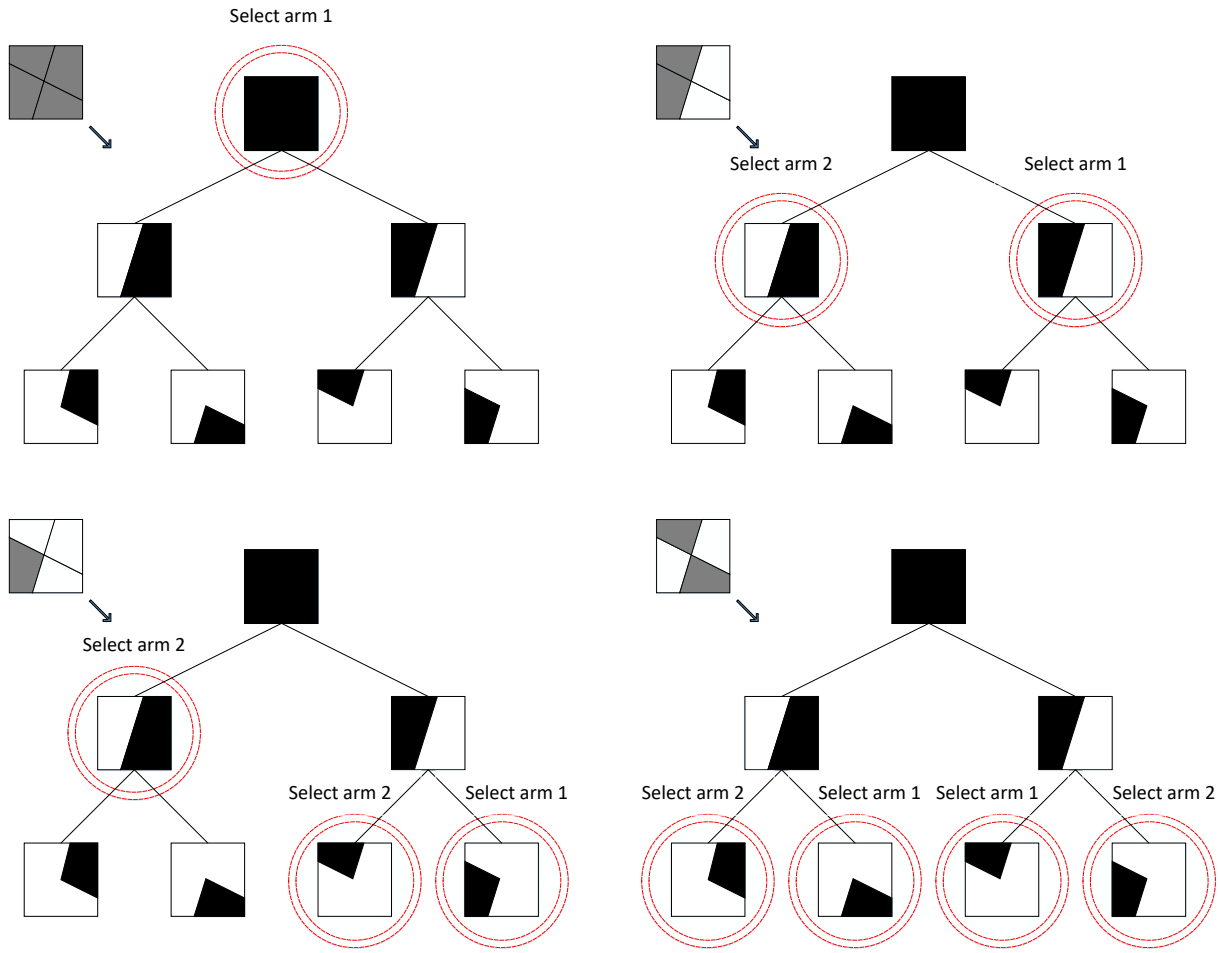


Fig. 4: Representation of 4 sample mappings in Fig. 2 over the binary tree in Fig. 3.

## V. A Weighted Mixture of Experts Algorithm Based on the Binary Tree

In this section, we use the binary tree defined in Section IV to implement the mixture of experts algorithm described in Section III. To this end, using a binary tree of depth $D$, we quantize the context space $S$ into $N = 2^D$ disjoint regions. Given that we have $M$ bandit arms, we define $M^N$ different experts corresponding to different assignment of arms to different quantized regions. Each expert is composed of a partition of the context space and an arm assigned to each region of this partition. As an example, consider a 2-armed bandit problem. Suppose that we use a binary tree of depth 2 to quantize the context space into 4 regions. In this case, we define $2^4 = 16$ experts as in Fig. 2. We have represented 4 samples among these 16 experts on our binary tree in Fig. 4. We seek to adaptively combine the defined experts to achieve the performance of the best one as explained in Section III. In order to implement our mixture of experts, over each node of the tree $v^{\{i,j\}}$ we define $M$ parameters $\alpha_{t,m}^{\{i,j\}}$ for $m = 1$ to $M$, as the weight of $m^{\text{th}}$ arm in the node $v^{\{i,j\}}$. This weight shows our trust on the $m^{\text{th}}$ arm when the context vector falls into the region corresponding to the node $v^{\{i,j\}}$. We set $\alpha_{1,m}^{\{i,j\}} = 1$ for all $m$'s and $v^{\{i,j\}}$'s, and for $t \geq 2$,

$$\alpha_{t,m}^{\{i,j\}} = \exp\left(-\eta \sum_{\tau=1}^{t-1} \frac{l_{I_\tau}}{p_{\tau,m}} \mathbf{1}_{\{I_\tau = m\}} \mathbf{1}_{\{\mathbf{s}_\tau \in R^{\{i,j\}}\}}\right). \quad (7)$$

Note that we can easily update these weight variables as follows. At each round $t$, after we observed $\mathbf{s}_t$, calculated $\mathbf{p}_t$, selected $I_t^{\text{th}}$ arm and suffered the loss $l_{t,I_t}$, we calculate

$$\alpha_{t+1,m}^{\{i,j\}} = \alpha_{t,m}^{\{i,j\}} \exp\left(-\eta \frac{l_{I_t}}{p_{t,m}} \mathbf{1}_{\{I_t = m\}} \mathbf{1}_{\{\mathbf{s}_t \in R^{\{i,j\}}\}}\right). \quad (8)$$

We point out that the weight of each of our $M^N$ experts in (4), i.e., $\alpha_{t,i}$, can be written as a multiplication of its initial weight and several weight parameters on the tree nodes, i.e., $\alpha_{t,m}^{\{i,j\}}$, corresponding to the mapping followed by the expert. As an example, as shown in Fig. 4, using the binary tree in Fig. 3 to represent the experts in Fig. 2, we have

$$\begin{aligned}
\alpha_{t,1} &= \alpha_{1,1} \times \left(\alpha_{t,1}^{\{0,1\}}\right), \\
\alpha_{t,6} &= \alpha_{1,6} \times \left(\alpha_{t,2}^{\{1,1\}} \alpha_{t,1}^{\{1,2\}}\right), \\
\alpha_{t,10} &= \alpha_{1,10} \times \left(\alpha_{t,2}^{\{1,1\}} \alpha_{t,2}^{\{2,3\}} \alpha_{t,1}^{\{2,4\}}\right), \\
\alpha_{t,10} &= \alpha_{1,10} \times \left(\alpha_{t,2}^{\{2,1\}} \alpha_{t,1}^{\{2,2\}} \alpha_{t,1}^{\{2,3\}} \alpha_{t,2}^{\{2,4\}}\right).
\end{aligned}$$

We define another variable $w_t^{\{i,j\}}$ over each node $v^{\{i,j\}}$ such that at the leaf nodes, where $i = D$,

$$w_t^{\{i,j\}} = \frac{1}{2M} \sum_{m=1}^{M} \alpha_{t,m}^{\{i,j\}}, \tag{9}$$

and at the lower layers, for $i = D-1$ to 0,

$$w_t^{\{i,j\}} = \frac{1}{2M} \sum_{m=1}^{M} \alpha_{t,m}^{\{i,j\}} + \frac{1}{2} w_t^{\{i+1,2j-1\}} w_t^{\{i+1,2j\}} - \frac{1}{8M^2} \sum_{m=1}^{M} \alpha_{t,m}^{\{i+1,2j-1\}} \alpha_{t,m}^{\{i+1,2j\}}. \tag{10}$$

We next show that using this recursion to calculate $w_t^{\{i,j\}}$ variables, $w_t^{\{0,1\}}$ is equal to $\sum_{k=1}^{M^N} \alpha_{t,k}$, where the initial weights of the experts, i.e., $\alpha_{1,i}$'s, are as follows. The initial weight of each of the $M$ experts which can be represented using one node of the tree (in 0$^{\text{th}}$ layer) are equal to $\frac{1}{2M}$. As the number of tree nodes needed to represent the experts increases by one, the initial weight of the experts gets multiplied by a factor of $\frac{1}{4M}$. Following our previous example, as shown in Fig. 4, the 1$^{\text{st}}$ expert has can be shown using one node, hence its initial weight $\alpha_{1,1}$ is equal to $\frac{1}{4}$; the 6$^{\text{th}}$ expert can be defined over 2 nodes, hence its initial weight $\alpha_{1,6}$ is equal to $\frac{1}{4} \times \frac{1}{8} = \frac{1}{32}$; the 14$^{\text{th}}$ expert is defined over 3 nodes, hence its initial weight $\alpha_{1,14}$ is equal to $\frac{1}{32} \times \frac{1}{8} = \frac{1}{256}$; and the 7$^{\text{th}}$ expert can be represented over 4 nodes, hence its initial weight $\alpha_{1,7}$ is equal to $\frac{1}{256} \times \frac{1}{8} = \frac{1}{2048}$.

Now, in order to calculate the probability simplex in (3), we define $M$ other variables to calculate $\sum_{j=1}^{M^N} \alpha_{t,j} \mathbf{1}_{\{g_j(\mathbf{s}_t)=i\}}$, for $i = 1, ..., M$, with the initial weights, i.e., $\alpha_{1,i}$'s, as described in the previous paragraph. To this end, after we observed $\mathbf{s}_t$, we set

$$\gamma_{t,m}^{\{D,j\}} = \frac{1}{2M} \alpha_{t,m}^{\{D,j\}}, \tag{11}$$

at the leaf node containing $\mathbf{s}_t$. Then we go up on the tree using a recursive formula similar to the way we calculate $w_t$ variables in (10). If we are going up on the left branch,

$$\gamma_{t,m}^{\{i,j\}} = \frac{1}{2M} \alpha_{t,m}^{\{i,j\}} + \frac{1}{2} \gamma_{t,m}^{\{i+1,2j-1\}} w_t^{\{i+1,2j\}} - \frac{1}{8M^2} \alpha_{t,m}^{\{i+1,2j\}} \alpha_{t,m}^{\{i+1,2j-1\}}, \tag{12}$$

and if we are going up on the right branch,

$$\gamma_{t,m}^{\{i,j\}} = \frac{1}{2M} \alpha_{t,m}^{\{i,j\}} + \frac{1}{2} \gamma_{t,m}^{\{i+1,2j\}} w_t^{\{i+1,2j-1\}} - \frac{1}{8M^2} \alpha_{t,m}^{\{i+1,2j\}} \alpha_{t,m}^{\{i+1,2j-1\}}. \tag{13}$$

Using this recursion we calculate $\gamma_{t,m}^{\{0,1\}}$ for $m = 1, ..., M$. We point out that $\gamma_{t,m}^{\{0,1\}}$ is the weighted sum of all experts, which select the $m^{\text{th}}$ arm when they observe $\mathbf{s}_t$. Hence, $\gamma_{t,m}^{\{0,1\}}$ is equal to $\sum_{j=1}^{M^N} \alpha_{t,j} \mathbf{1}_{\{g_j(\mathbf{s}_t)=m\}}$ for different $m$'s. Therefore, we can build the probability simplex in (3) as follows

$$p_{t,i} = \gamma_{t,m}^{\{0,1\}} / w_t^{\{0,1\}}, \forall i \in 1, ..., M. \tag{14}$$

With the proposed implementation of the algorithm, at each round after observing $\mathbf{s}_t$, we first calculate $\gamma_{t,m}^{\{0,1\}}$ for $m = 1, ..., M$. Calculating each of these variables needs $D = \log_2 N$ recursion, i.e., the number of tree layers. Then, we divide the $\gamma$ variables by $w_t^{\{0,1\}}$ to form the probability simplex $\mathbf{p}_t = (p_{t,1}, ..., p_{t,M})$, using which we select an arm $I_t$. After we selected our arm and suffered the loss according to the selected arm, we first update $\alpha_{t,I_t}^{\{i,j\}}$ parameters at the nodes containing $\mathbf{s}_t$. Then, we update $w_t^{\{i,j\}}$ variables at these affected nodes and go to the next round. Therefore, the computational complexity at each round is of $O(M \ln N)$. On the other hand, a binary tree of depth $D = \log_2 N$ has just $2N - 1$ nodes and we store $M + 1$ variables, i.e., $\{\alpha_{t,1}^{\{i,j\}}, ..., \alpha_{t,M}^{\{i,j\}}, w_t^{\{i,j\}}\}$, for each node. Hence, the space complexity of our algorithm is of $O(MN)$. The pseudo code of explained procedure is provided in Algorithm 1.

In the sequel we derive an upper bound of $O(\sqrt{MTR \ln N \ln M})$ for the regret of our algorithm against the best best mapping in $C^N$, given that the optimal mapping has $R$ disjoint regions. To this end, we start by showing how the normalized weight of each expert in our algorithm is lower bounded.

**Lemma 1.** *If we assign an initial weight equal to* $(1/2M) \times (1/4M)^{R_b - 1}$ *to each expert with* $R_b$ *nodes, denoting the initial weight of* $i^{th}$ *expert by* $\alpha_{1,i}$*, we have*

$$\sum_{i=1}^{M^N} \alpha_{1,i} \leq 1. \tag{15}$$

*Hence, the normalized initial weight of each expert with* $R_b$ *nodes is lower bounded as follows*

$$\beta_{1,i} \geq \frac{1}{(2M)(4M)^{R_b - 1}} \geq \frac{1}{(4M)^{R_b}}. \tag{16}$$

*Proof of Lemma 1:* From the definition $w_1^{\{0,1\}} = \sum_{k=1}^{M^N} \alpha_{t,k}$, therefore, it suffices to upper bound $w_1^{\{0,1\}}$. Since we set $\alpha_{1,m}^{\{i,j\}} = 1$ for all nodes, at the leaf nodes, where $i = D$, from (9) we have $w_1^{\{i,j\}} = 0.5$. As we go up on the tree, for $i = D - 1$ to 0, from (10) we have

$$w_1^{\{i,j\}} = \frac{1}{2} + \frac{1}{2} w_1^{\{i+1,2j-1\}} w_1^{\{i+1,2j\}} - \frac{1}{8M} \leq \frac{1}{2} + \frac{1}{2} w_1^{\{i+1,2j-1\}} w_1^{\{i+1,2j\}}. \tag{17}$$

Hence, as long as $w_1^{\{i,j\}}$ is less than or equal to 1 for all nodes in the $i^{\text{th}}$ layer of the tree, it remains less than or equal to 1 in the $(i-1)^{\text{th}}$ level, and the claim of the Lemma holds by induction. ∎

Now, in the following lemma we show the relationship between the number of disjoint regions of an expert and the number of its nodes.

**Lemma 2.** *Denoting the number of an expert's nodes by* $R_b$*, and the number of disjoint regions of the expert by* $R$*, the following bound holds*

$$R_b \leq 1 + (R - 1) \log_2 N. \tag{18}$$

*Proof:* If $R_b$ is equal to 1, $R$ is also equal to 1. As $R$ increases by one, at the worst case scenario we have to split

**Algorithm 1** Binary tree based Exponential weights for Bandits with Known number of disjoint regions (BEB.K)

**Parameter:**
Set constant $\eta \in \mathbb{R}^+$
**Initialization:**
Initialize $\alpha_{1,m}^{\{i,j\}} = 1$ for all $m, \{i,j\}$
Initialize $w_1^{\{i,j\}}$ for all $\{i,j\}$ using (9) and (10)
**Algorithm:**
**for** $t = 1$ **to** $T$ **do**
  Observe $\mathbf{s}_t \in \{1, 2, ..., 2^D\}$
  **for** $m = 1$ **to** $M$ **do**
    $i = D$ , $j = \mathbf{s}_t$
    $x = 1$ iff $j$ is odd and $x = 0$ otherwise
    $\gamma_{t,m}^{\{i,j\}} = (1/2M)\alpha_{t,m}^{\{i,j\}}$
    **while** $i \geq 1$ **do**
      $i = i - 1$
      $j = \lceil j/2 \rceil$
      if $x = 1$
      Calculate $\gamma_{t,m}^{\{i,j\}}$ using (12)
      if $x = 0$
      Calculate $\gamma_{t,m}^{\{i,j\}}$ using (13)
      $x = 1$ iff $j$ is odd and $x = 0$ otherwise
    **end while**
  **end for**
  **for** $m = 1$ **to** $M$ **do**
    $p_{t,m} = \gamma_{t,m}^{\{0,1\}} / w_t^{\{0,1\}}$
  **end for**
  Select a random arm $I_t$ according to $\mathbf{p}_t = (p_{t,1}, ..., p_{t,M})$
  $\alpha_{t+1,m}^{\{i,j\}} = \alpha_{t,m}^{\{i,j\}}$ for all $m, \{i,j\}$
  $w_{t+1}^{\{i,j\}} = w_t^{\{i,j\}}$ for all $\{i,j\}$
  $i = D$ , $j = \mathbf{s}_t$
  **while** $i \geq 0$ **do**
    $\alpha_{t+1,I_t}^{\{i,j\}} = \alpha_{t,I_t}^{\{i,j\}} \exp\left(-\eta \frac{l_{t,I_t}}{p_{t,I_t}}\right)$
    $i = i - 1$
    $j = \lceil j/2 \rceil$
  **end while**
  $i = D, j = \mathbf{s}_t$
  Calculate $w_{t+1}^{\{i,j\}}$ using (9)
  **while** $i \geq 1$ **do**
    $i = i - 1$
    $j = \lceil j/2 \rceil$
    Calculate $w_{t+1}^{\{i,j\}}$ using (10)
  **end while**
**end for**

the context space for $D = \log_2 N$ times to build the new border, which has caused $R$ to increase. Hence, $R_b$ increases by $D$ at most and the bound holds. ∎

Putting (16) and (18) into (6) directly concludes the following theorem.

**Theorem 2.** *Given the number of disjoint regions of the optimal expert $R$, the regret of Algorithm 1 against any mapping in $C^N$ satisfies*

$$\mathcal{R}(T, C^N) \leq \frac{R \log_2 N \ln 4M}{\eta} + \frac{MT\eta}{2}. \tag{19}$$

**Corollary 1.** *In the case of known number of rounds $T$, choosing $\eta = \sqrt{\frac{2R \log_2 N \ln 4M}{MT}}$ directly yields*

$$\mathcal{R}(T, C^N) \leq \sqrt{2RMT \log_2 N \ln 4M}. \tag{20}$$

We have successfully achieved a regret bound of $O(\sqrt{MTR \ln N \ln M})$ for the case when $R$ and $T$ are known beforehand. In the following section, we extend our algorithm to achieve a regret bound of the same order without any a priori information on these parameters.

## VI. Extension to the Case of Unknown Parameters

In this section, we first introduce an algorithm that achieve a regret upper bound of $O(\sqrt{MTR \ln N \ln M})$ in the case of unknown $R$ and known $T$. Then, we explain how our algorithms can get straightforwardly extended to the case of unknown $T$ as well.

### A. Extension to the Case of Unknown Number of Disjoint Regions

In the previous section, we introduced Algorithm 1 that achieves the regret bound of $O(\sqrt{MTR \ln N \ln M})$ given the parameters $R$ and $T$ a priori, where Algorithm 1 needs these parameters to tune its $\eta$ parameter. We seek to remove the requirement of knowing $R$, while keeping the computational complexity logarithmically dependent on the number of quantization levels $N$.

To this end, we run $D = \log_2 N$ copies of Algorithm 1 in parallel, where the $\eta$ parameters in different parallel algorithms are optimized for different $R$ parameters $\tilde{R} \in \{1, 2, 4, ..., 2^D\}$. Note that one of the elements of the set of $\tilde{R}$'s is such that $\tilde{R} \leq R \leq 2\tilde{R} - 1$. We consider each of these $D = \log_2 N$ parallel algorithms as an expert $\mathcal{E}_{\tilde{R}}$ and combine them using an *EXP4* algorithm to achieve the performance of the best algorithm in the mixture. The detailed pseudo code of this combination is provided in Algorithm 2. The following theorem shows a regret bound of $O(\sqrt{MTR \ln N \ln M})$ for Algorithm 2 against the optimal mapping from the quantized regions to the arms set.

**Theorem 3.** *Given the number of quantization levels equal to $N = 2^D$, the regret of Algorithm 2, i.e., BEB.UK, against the class of all mapping from the quantized regions to the bandit arms, i.e., $C^N$, is upper bounded as follows*

$$R(T, C^N) \leq \frac{3}{2}\sqrt{MTD(R+1) \ln 4M} + \sqrt{2MT \ln D}, \tag{21}$$

Note that the parameter $D = \log_2 N$, hence Algorithm 2 achieves a regret bound of $O(\sqrt{MTR \ln N \ln M})$ without any a priori information on the parameter $R$ of the optimal mapping.

**Algorithm 2** Binary tree based Exponential weights for Bandits with Unknown number of disjoint regions (BEB.U)

**Initialization:**
Initialize $D$ experts $\mathcal{E}_{\tilde{R}}$ for $\tilde{R} = 2^i$, $i \in \{1, 2, ..., D\}$. Each expert $\mathcal{E}_{\tilde{R}}$ uses *BEB.K* with $\eta$ parameter tuned for $\tilde{R}$ disjoint regions and suggests the probability simplex $\mathbf{p}_t^{\tilde{R}}$ to select a bandit arm at round $t$.
Initialize expert probability vector $\mathbf{q}_1 = (q_{1,1}, ..., q_{1,D})$ to be uniform
**Algorithm:**
**for** $t = 1$ **to** $T$ **do**
  Observe $\mathbf{s}_t \in \{1, 2, ..., 2^D\}$
  Select one of the $D$ experts according to the probability simplex $\mathbf{q}_t$
  Select an arm according to the suggestion of the selected expert $\mathbf{p}_t^{\tilde{R}}$ given $\mathbf{s}_t$.
  Receive loss of the selected arm $l_{t,I_t}$.
  Update $\mathbf{q}_t$ according to *EXP4* algorithm
  Feed the observed context vector $\mathbf{s}_t$, the selected arm $I_t$ and the received loss $l_{t,I_t}$ to all $D$ experts and update them using *BEB.K* algorithm.
**end for**

*Proof of Theorem 3:* Denoting the arm chosen by the expert $\mathcal{E}_{\tilde{R}}$ at round $t$ by $I_t^{\tilde{R}}$, using *EXP4* algorithm [8] yields

$$\mathbb{E}[\sum_{t=1}^{T} l_{t,I_t} - \sum_{t=1}^{T} l_{t,I_t^{\tilde{R}}}] \leq \sqrt{2MT \ln D}. \tag{22}$$

On the other hand, one of our $D$ experts is such that $\tilde{R} \leq R \leq 2\tilde{R} - 1$. For this expert, from (19) we have

$$\max_{g \in C^N} \mathbb{E}[\sum_{t=1}^{T} l_{t,I_t^{\tilde{R}}} - \sum_{t=1}^{T} l_{t,g(\mathbf{s}_t)}] \leq \frac{RD \ln 4M}{\eta} + \frac{MT\eta}{2}, \tag{23}$$

where

$$\eta = \sqrt{\frac{2\tilde{R}D \ln 4M}{MT}}, \tag{24}$$

since this expert sets its learning parameter $\eta$ assuming that the optimal strategy has $\tilde{R}$ disjoint regions. Putting this $\eta$ into (23), and combining the result with (22) leads to the following regret bound

$$R(T, C^N) \leq \sqrt{2MT \ln D} + (\frac{R + \tilde{R}}{\sqrt{\tilde{R}}}) \sqrt{\frac{MTD \ln 4M}{2}}. \tag{25}$$

It can be easily shown that if $\tilde{R} \leq R \leq 2\tilde{R} - 1$, we have

$$\frac{R + \tilde{R}}{\sqrt{\tilde{R}}} \leq \frac{3R + 1}{\sqrt{2(R + 1)}} < 3\sqrt{\frac{R + 1}{2}}. \tag{26}$$

Putting this upper bound in (25) concludes the proof. ∎

We emphasize that the space and computational complexities of Algorithm 2 is simply $\log_2 N$ times those of Algorithm 1. Therefore, Algorithm 2, while achieving a regret upper bound of the same order as Algorithm 1, removes the requirement of knowing the parameter $R$, with cost of an increment in the space and computational complexities by a multiplicative factor of $\log_2 N$.

Now we extend our algorithm to achieve the regret bound of $O(\sqrt{MTR \ln N \ln M})$, when we have no knowledge of the parameter $T$ as well.

**Algorithm 3** BEB.U with Doubling trick (BEB.U.D)

**for** $n = 0, 1, 2, ...$ **do**
  Run Algorithm 2 with $\eta$ parameters tuned for a $2^n$ rounds game on rounds $t = 2^n, ..., 2^{n+1} - 1$
**end for**

### *B. Extension to the Case of Unknown Game Length*

In the case when we do not know the game length $T$ to use for setting the parameters of our algorithm, we can use the doubling trick [37] and run Algorithm 3. It is shown in [37] that if the regret bound of an algorithm is of the form $x\sqrt{T}$, the doubling procedure used in Algorithm 3 yields the regret upper bound of $\frac{\sqrt{2}}{\sqrt{2}-1} x\sqrt{T}$. Therefore, we get the following result when we do not know the game length beforehand.

**Corollary 2.** *If $R$ and $T$ are not known a priori, then Algorithm 3, i.e., BEB.U, has the regret bound*

$$R(T, C^N) \leq \frac{3}{2 - \sqrt{2}} \sqrt{MTD(R + 1) \ln 4M} + \frac{2}{\sqrt{2} - 1} \sqrt{MT \ln D}. \tag{27}$$

Therefore, in the case of unknown number of rounds $T$ Algorithm 3 achieves a regret bound of the same order as Algorithm 2 with cost of a multiplicative constant factor in the regret upper bound.

In the following, we introduce an efficient method for quantizing the context space. Using this method, in the case when the arm loss functions are Lipschitz continuous our algorithm asymptotically achieves a performance as close as desired to the performance of the truly optimal mapping by increasing the number of quantization levels $N$.

## VII. An Efficient Quantization Method to Asymptotically Achieve the Optimal Context Based Arm Selection

Suppose that the context space is the $n$ dimensional space $S = [0, 1]^n$. Given the binary tree's depth $D$, our quantization scheme is as follows. We split the context space into $2^{\lfloor D/n \rfloor + 1}$ equal subspaces along the first $D \pmod n$ dimensions (of the total $n$ dimensions), and $2^{\lfloor D/n \rfloor}$ equal subspaces along the remaining dimensions.

**Theorem 4.** *Using aforementioned quantization method for our algorithm, if the arm loss functions are Lipschitz continuous with the Lipschitzness constant equal to c, the difference between the loss corresponding to the best mapping in $C^N$ and the loss corresponding to the truly optimal mapping (in $C$) is upper bounded by*

$$\frac{2c\sqrt{n}}{\sqrt[n]{N}} \tag{28}$$

*Proof of Theorem 3:* Using this quantization method, the subspaces in the finest partition of the context space are $n$-dimensional cubes with the longest diagonal length equal to

$$\sqrt{\frac{n-(D\ (\mathrm{mod}\ n))}{(2^{\lfloor \frac{D}{n} \rfloor})^2}+\frac{D\ (\mathrm{mod}\ n)}{(2^{\lfloor \frac{D}{n}+1 \rfloor})^2}}. \tag{29}$$

Since $D\ (\mathrm{mod}\ n) \geq 0$, this upper bound is at most equal to

$$\sqrt{\frac{n}{2^{2\lfloor \frac{D}{n} \rfloor}}} \leq \frac{2\sqrt{n}}{2^{\frac{D}{n}}} = \frac{2\sqrt{n}}{\sqrt[n]{N}}. \tag{30}$$

Since the loss functions are Lipschitz continuous, the difference between the loss corresponding to the truly optimal mapping and the best mapping in $C^N$ cannot exceed the Lipschitzness constant times the quantized cubes diagonal length, which concludes the proof. ■

Note that the Lipschitzness assumption does not intervene with the adversarial setting. The loss functions can be quite different in different rounds and as long as they are Lipschitz continuous at each specific round, the assumption holds and our algorithm is competitive against the class of all possible mapping $C$. In this case, combining (27) and (28) directly concludes the following theorem.

**Theorem 5.** *Consider a contextual $M$-armed bandit problem with the context space $S = [0,1]^n$, where the loss functions of the arms are Lipschitz continuous with the constant $c$ at all rounds. If we quantize the context space into $N = 2^D$ disjoint regions using the quantization scheme described in Section VII, the regret of Algorithm 1 against the truly optimal strategy in a $T$ round trial is upper bounded as follows*

$$R(T,C) \leq \frac{3}{2-\sqrt{2}}\sqrt{MT\log_2 N(R+1)\ln 4M} + \frac{2}{\sqrt{2}-1}\sqrt{MT\ln(\log_2 N)} + \frac{2Tc\sqrt{n}}{\sqrt[n]{N}}. \tag{31}$$

We emphasize that we can make the linear-in-time term of the upper bound in (31) as small as desired by increasing the depth parameter of our algorithm $D = \log_2 N$ (or equivalently increasing the quantization level $N$).

## VIII. Experiments

In this section, we demonstrate the performance of our algorithm in different scenarios involving both real and synthetic data. We compare the performance of our algorithms *BEB.K* and *BEB.U* with various depth parameters against the state-of-the-art adversarial bandit algorithms *EXP3* and *S-EXP3* [8]. In all of the experiments, the parameters of *EXP3* and *S-EXP3* algorithms are set to their optimal values according to their publication [8].

### A. Stationary Environment

We first construct a game with 3-armed bandit, where the context space is the 1-dimensional space $S = [0,1]$. Each arm $i$, generates its loss according to a Bernoulli distribution with parameter $p_i$, i.e., the loss is equal to 1 with probability equal to $p_i$. These parameters, i.e., $p_1, p_2, p_3$, depend on the context variable $s_t$ as

$$\begin{aligned} p_1(s_t) &= 0.5 + 0.5\sin(2\pi s_t), \\ p_2(s_t) &= \sin(\pi s_t), \\ p_3(s_t) &= s_t. \end{aligned} \tag{32}$$

Note that here the optimal strategy is defined as follows

$$g(s_t) = \begin{cases} 3 & , s_t < 0.5 \\ 1 & , 0.5 \leq s_t < 0.9182 \\ 2 & , 0.9182 \leq s_t. \end{cases} \tag{33}$$

In this experiment, we generated the context variable $s_t$ randomly with uniform distribution over the context space, i.e., $[0,1]$, and compared the averaged cumulated loss performance, i.e., $(\sum_{\tau=1}^{t} l_{\tau,I_\tau})/t$, for our algorithm *BEB.U* with various depth parameters equal to 2, 5, and 10, *S-EXP3* [8] with the same depth parameters, and *EXP3* [8].

To this end, we generated 10 synthetic datasets of length $10^5$. To produce each dataset, first, $10^5$ context variables $s_t$ are drawn according to a uniform probability distribution over the interval $[0,1]$. Then, the arm losses corresponding to different rounds are drawn from the Bernoulli distributions, parameters of which are determined according to (32). Each dataset is presented to the algorithms 10 times and the results are averaged. This process is repeated for all 10 datasets and the ensemble averages are plotted in Fig. 5. Two important results can be derived from this figure. First, our algorithm *BEB.U* with a specific depth outperforms *S-EXP3* with the same depth, as well as *EXP3* algorithm in this scenario. Second, while increment in the depth parameter uniformly

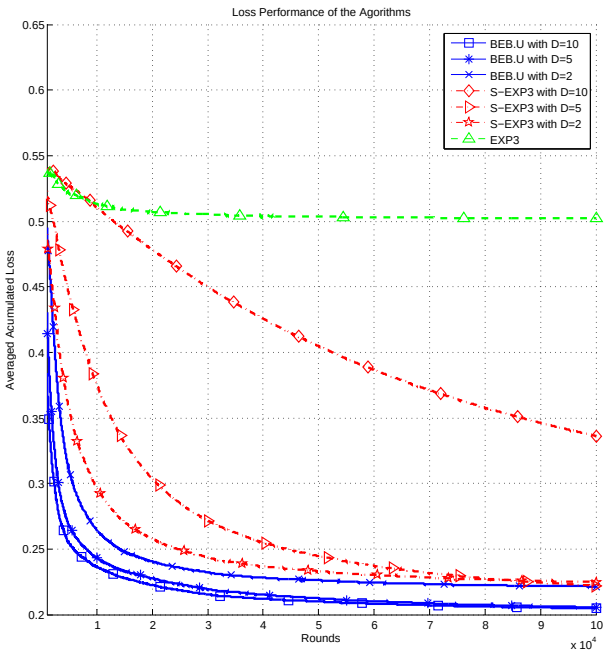


Fig. 5: The averaged accumulated loss of *BEB.U*, *S-EXP3*, and *EXP3*, averaged on 10 synthetic datasets defined using (32), where the result for each dataset is averaged over 10 different trials.

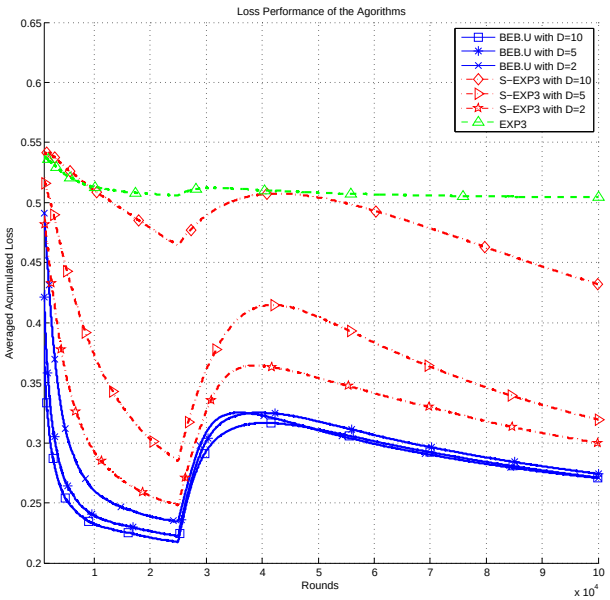


Fig. 6: The averaged accumulated loss of *BEB.U*, *S-EXP3*, and *EXP3*, averaged on 10 synthetic datasets as described in Section VIII-B, involving a rapid change in the behavior of the arms after 25% of the rounds. The result for each data set is averaged over 10 different trials.

improves the performance of our algorithm, it can degrade the performance of *S-EXP3* due to the overtraining problem. The superior performance our algorithm in this experiment is because of its fast convergence to the optimal mapping. Here, *EXP3* has a fast convergence but it converges to a suboptimal mapping because it does not use the context information. On the other hand, *S-EXP3* converges to the optimal mapping, but needs a huge amount of data to get trained. Our algorithm uses an efficient adaptive combination of the experts with intelligent initial weights to obtain the advantages of both *EXP3* and *S-EXP3* algorithms, while mitigating their disadvantages. We emphasize that even though the $R$ parameter was known to be equal to 3 from (33), we did not use this information in favor of ourself and used *BEB.U* algorithm, which does not need to know $R$ in order to optimize its parameters.

### B. Nonstationary Environment

In this part, we illustrate the averaged cumulated loss performance of the algorithms in a nonstationary environment. To this end, we construct 10 different datasets of length $10^5$ as in Section VIII-A. However, here the arm losses follow a model as in (32) in the first quarter of the rounds, and the following model in the rest of the rounds:

$$\begin{aligned} p_1(s_t) &= \sin(\pi s_t), \\ p_2(s_t) &= s_t, \\ p_3(s_t) &= 0.5 + 0.5\sin(2\pi s_t). \end{aligned} \tag{34}$$

Hence, we have an abrupt change in the model of the arms within the rounds. Each dataset is presented to the algorithms 10 times and the results are averaged. This process is repeated for all 10 datasets and the ensemble averages are plotted in Fig. 6. As shown in the figure, our algorithm *BEB.U*, not only outperforms its competitor before the rapid change in the model of the bandit arms, but also adopts better to this rapid change in comparison to the competitors. Note that in order to have a fair comparison, we used the *BEB.U* version of our algorithm which requires no information about $R$ to optimize its parameters.

**Algorithm 4** The offline evaluation method used to test the competitor algorithms over the Yahoo! Today Module dataset

**Input:** Bandit algorithm $\mathcal{A}$, logged data for $T$ rounds
**Initialize:** $L = 0$ and $R = 0$
**for** $t = 1$ **to** $T$ **do**
  Get $s_t \in \{1, 2, ..., N\}$ from the log
  Run the algorithm $\mathcal{A}$.
  **if** the arm, selected by $\mathcal{A}$ is the arm which is shown to the user **then**
    Use the user feedback to update $\mathcal{A}$.
    Set $R = R + 1$.
    If the user has not clicked set $L = L + 1$.
  **else**
    Ignore this round.
  **end if**
**end for**
$L$ and $R$ show the total loss and the total rounds respectively.

### C. Real Dataset

In this section, we demonstrate the superior performance of our algorithms *BEB.K* and *BEB.U* against their natural competitors *EXP3* and *S-EXP3* over the well known real life dataset provided by Yahoo! Research. This dataset contains a user click log for news articles displayed in the featured tab of the Today Module on Yahoo!'s front page, within October 2 to 16, 2011. The dataset contains 28041015 user visits. For each visit, the user is associated with a binary feature vector of dimension 136 that contains information about the user like age, gender, behavior targeting features, etc. We used an unbiased offline evaluation method as in [38], to test the competitors over this dataset. A brief pseudo-code of this evaluation method is shown in Algorithm 4. In this experiment, we ran a PCA algorithm [39] over the first 5% of the data to get the principal components of the feature vectors. We mapped the feature vectors over the first principal component to form a set of $1-$dimensional context variables. We used these context variables for *S-EXP3*, *BEB.K* and *BEB.U* algorithms. We tested the *EXP3* and *S-EXP3* algorithms with several depth parameters, while their parameters were set to their optimum values [8]. The parameters of *BEB.U* are also set to their optimal values. However, since we do not have any information about the number of disjoint regions in the optimal mapping, i.e., $R$, the $\eta$ parameter for the *BEB.K* algorithm can not be tuned to the optimum value analytically. In this experiment, in order to have a fair comparison, we set the $\eta$ parameter of the *BEB.K* algorithm with a specific depth parameter, equal to the $\eta$ parameter of the *S-EXP3* algorithm with the same depth.

TABLE I: Percentage of click in the Yahoo! Today Module dataset

| Algorithms \ Depth | 2 | 5 | 10 |
|---|---|---|---|
| Random | 3.68 | 3.68 | 3.68 |
| EXP3 | 3.69 | 3.69 | 3.69 |
| S-EXP3 | 3.70 | 3.70 | 3.73 |
| BEB.K | 4.48 | 4.57 | 4.63 |
| BEB.U | 4.64 | 4.63 | 4.63 |

We emphasize that no numerical optimization is done for the $\eta$ parameter of *BEB.K* algorithm. The percentage of user clicks for different algorithms are shown in Table 1. As shown in this table, our algorithm *BEB.K* with a specific depth outperforms *S-EXP3* with the same depth as well as *EXP3* algorithm, even though its learning rate parameter is not tuned to the optimum value due to the lack of knowledge about the parameter $R$. Moreover, due to the unknown $R$, *BEB.U* algorithm has a superior performance in comparison to *BEB.K*, as expected.

## IX. Concluding Remarks

We studied the contextual multi-armed bandit problem in an adversarial setting and introduced truly online and low complexity algorithms that asymptotically achieve the performance of the best context dependent bandit arm selection policy. Our core algorithm quantizes the space of the context vectors into a large number of disjoint regions using an efficient quantization method and forms the class of all mappings from these regions to the bandit arms. Then, it adaptively combines these mappings in a mixture-of-experts setting and achieves the performance of the best mapping in the class. We run a specific number of copies of our core algorithm with different parameters and combine them to achieve the performance of the best one with the optimal parameter. We proved performance upper bounds for the introduced algorithms. These upper bounds show that we achieve a performance as close as desired to the performance of the truly optimal policy, by increasing the number of quantization levels. We used a novel notion of a binary tree to implement our algorithms in an efficient way such that the computational complexity is log-linear in the number of quantization levels. We have no statistical assumptions on the behavior of the context vectors and the bandit arms, hence our results are guaranteed to hold in an individual sequence manner. Through extensive set of experiments involving synthetic and real data, we demonstrated the significant performance gains achieved by the proposed algorithms in comparison to the state-of-the-art adversarial multi-armed bandit algorithms.

## Appendix
## Proof of Theorem 1

From the definition, denoting the mapping followed by the $j^{\text{th}}$ expert by $g_j(.)$, we have

$$\mathcal{R}(T,E_j)=\mathbb{E}\left[\sum_{t=1}^{T} l_{t,I_t}-\sum_{t=1}^{T} l_{t,g_j(\mathbf{S}_t)}\right] \tag{35}$$

Here, $l_{t,I_t}$ can be expanded as

$$\begin{aligned} l_{t,I_t} &= \mathbb{E}_{j\sim\boldsymbol{\beta}_t}\tilde{l}_{t,g_j(\mathbf{S}_t)} \\ &= \frac{1}{\eta}\left(\ln\left(\mathbb{E}_{j\sim\boldsymbol{\beta}_t}e^{-\eta\tilde{l}_{t,g_j(\mathbf{S}_t)}}\right)+\eta\mathbb{E}_{j\sim\boldsymbol{\beta}_t}\tilde{l}_{t,g_j(\mathbf{S}_t)}\right) \\ &\quad -\frac{1}{\eta}\ln\mathbb{E}_{j\sim\boldsymbol{\beta}_t}e^{-\eta\tilde{l}_{t,g_j(\mathbf{S}_t)}}. \end{aligned} \tag{36}$$

The first term in (36) can be bounded using the inequalities $\ln x\le x-1$ and $\exp(-x)-1+x\le x^2/2$, for all $x\ge 0$, as

$$\begin{aligned} &\ln\left(\mathbb{E}_{j\sim\boldsymbol{\beta}_t}e^{-\eta\tilde{l}_{t,g_j(\mathbf{S}_t)}}\right)+\eta\mathbb{E}_{j\sim\boldsymbol{\beta}_t}\tilde{l}_{t,g_j(\mathbf{S}_t)} \\ &\le \mathbb{E}_{j\sim\boldsymbol{\beta}_t}\left[e^{-\eta\tilde{l}_{t,g_j(\mathbf{S}_t)}}-1+\eta\tilde{l}_{t,g_j(\mathbf{S}_t)}\right] \\ &\le \mathbb{E}_{j\sim\boldsymbol{\beta}_t}\frac{\eta^2\tilde{l}^2_{t,g_j(\mathbf{S}_t)}}{2}=\frac{\eta^2 l^2_{t,I_t}}{2p_{t,I_t}}\le\frac{\eta^2}{2p_{t,I_t}}. \end{aligned} \tag{37}$$

In order to bound the second term in (36), we just rewrite the expectation using (4) as follows. For $t=1$, we have

$$-\frac{1}{\eta}\ln\mathbb{E}_{j\sim\boldsymbol{\beta}_1}e^{-\eta\tilde{l}_{1,g_j(\mathbf{S}_1)}}=-\frac{1}{\eta}\ln\frac{\sum_{j=1}^{M^N}\alpha_{1,j}e^{-\eta\tilde{l}_{1,g_j(\mathbf{S}_1)}}}{\sum_{j=1}^{M^N}\alpha_{1,j}}, \tag{38}$$

and for $t\ge 2$, we have

$$\begin{aligned} &-\frac{1}{\eta}\ln\mathbb{E}_{j\sim\boldsymbol{\beta}_t}e^{-\eta\tilde{l}_{t,g_j(\mathbf{S}_t)}} \\ &\qquad=-\frac{1}{\eta}\ln\frac{\sum_{j=1}^{M^N}\alpha_{1,j}e^{-\eta\sum_{\tau=1}^{t}\tilde{l}_{\tau,g_j(\mathbf{S}_\tau)}}}{\sum_{j=1}^{M^N}\alpha_{1,j}e^{-\eta\sum_{\tau=1}^{t-1}\tilde{l}_{\tau,g_j(\mathbf{S}_\tau)}}}. \end{aligned} \tag{39}$$

Putting the bounds in (37) and (39) into (36), we have

$$\begin{aligned} \sum_{t=1}^{T} l_{t,I_t} &\le -\frac{1}{\eta}\Bigg(\sum_{t=2}^{T}\ln\frac{\sum_{j=1}^{M^N}\alpha_{1,j}e^{-\eta\sum_{\tau=1}^{t}\tilde{l}_{\tau,g_j(\mathbf{S}_\tau)}}}{\sum_{j=1}^{M^N}\alpha_{1,j}e^{-\eta\sum_{\tau=1}^{t-1}\tilde{l}_{\tau,g_j(\mathbf{S}_\tau)}}} \\ &\quad+\ln\frac{\sum_{j=1}^{M^N}\alpha_{1,j}e^{-\eta\tilde{l}_{1,g_j(\mathbf{S}_1)}}}{\sum_{j=1}^{M^N}\alpha_{1,j}}\Bigg)+\frac{\eta T}{2p_{t,I_t}}. \end{aligned} \tag{40}$$

Opening the first two term in (40), we have

$$\begin{aligned} \sum_{t=1}^{T} l_{t,I_t} &\le -\frac{1}{\eta}\ln\sum_{j=1}^{M^N}\alpha_{1,j}e^{-\eta\sum_{\tau=1}^{T}\tilde{l}_{\tau,g_j(\mathbf{S}_\tau)}} \\ &\quad+\frac{1}{\eta}\ln\sum_{j=1}^{M^N}\alpha_{1,j}+\frac{\eta T}{2p_{t,I_t}}. \end{aligned} \tag{41}$$

Since $\sum_{j=1}^{M^N}\alpha_{1,j}e^{-\eta\sum_{\tau=1}^{T}\tilde{l}_{\tau,g_j(\mathbf{S}_\tau)}}\le\alpha_{1,j}e^{-\eta\sum_{\tau=1}^{T}\tilde{l}_{\tau,g_j(\mathbf{S}_\tau)}}$, we have

$$\begin{aligned} \sum_{t=1}^{T} l_{t,I_t} &\le -\frac{1}{\eta}\ln\alpha_{1,j}+\sum_{\tau=1}^{T}\tilde{l}_{\tau,g_j(\mathbf{S}_\tau)} \\ &\quad+\frac{1}{\eta}\ln\sum_{j=1}^{M^N}\alpha_{1,j}+\frac{\eta T}{2p_{t,I_t}} \\ &=\frac{\ln 1/\beta_{1,j}}{\eta}+\frac{\eta T}{2p_{t,I_t}}+\sum_{\tau=1}^{T}\tilde{l}_{\tau,g_j(\mathbf{S}_\tau)}. \end{aligned} \tag{42}$$

Taking expectation from both sides (with respect to $I_t\sim\mathbf{p}_t$) and substituting $\mathbb{E}[\tilde{l}_{\tau,g_j(\mathbf{S}_\tau)}]=l_{\tau,g_j(\mathbf{S}_\tau)}$ and $\mathbb{E}[\frac{1}{p_{t,I_t}}]=M$ into the result concludes the proof.

## REFERENCES


[1] Robert J. Meyer and Yong Shi, "Sequential choice under ambiguity: Intuitive solutions to the armed-bandit problem," *Management Science*, vol. 41, no. 5, pp. 817–834, 1995.

[2] Shai Shalev-Shwartz, "Online learning and online convex optimization," *Found. Trends Mach. Learn.*, vol. 4, no. 2, pp. 107–194, Feb. 2012.

[3] Nicol Cesa-Bianchi and Gbor Lugosi, "Combinatorial bandits," *Journal of Computer and System Sciences*, vol. 78, no. 5, pp. 1404 – 1422, 2012, {JCSS} Special Issue: Cloud Computing 2011.

[4] A. J. Bean and A. C. Singer, "Universal switching and side information portfolios under transaction costs using factor graphs," *IEEE Journal of Selected Topics in Signal Processing*, vol. 6, no. 4, pp. 351–365, Aug 2012.

[5] A. C. Singer, S. S. Kozat, and M. Feder, "Universal linear least squares prediction: upper and lower bounds," *IEEE Transactions on Information Theory*, vol. 48, no. 8, pp. 2354–2362, Aug 2002.

[6] A. C. Singer and M. Feder, "Universal linear prediction by model order weighting," *IEEE Transactions on Signal Processing*, vol. 47, no. 10, pp. 2685–2699, Oct 1999.

[7] T. Moon and T. Weissman, "Universal fir mmse filtering," *IEEE Transactions on Signal Processing*, vol. 57, no. 3, pp. 1068–1083, March 2009.

[8] Sébastien Bubeck and Nicolò Cesa-Bianchi, "Regret analysis of stochastic and nonstochastic multi-armed bandit problems," *CoRR*, vol. abs/1204.5721, 2012.

[9] Lihong Li, Wei Chu, John Langford, and Robert E. Schapire, "A contextual-bandit approach to personalized news article recommendation," in *Proceedings of the 19th International Conference on World Wide Web*, New York, NY, USA, 2010, WWW '10, pp. 661–670, ACM.

[10] C. Tekin, S. Zhang, and M. van der Schaar, "Distributed online learning in social recommender systems," *IEEE Journal of Selected Topics in Signal Processing*, vol. 8, no. 4, pp. 638–652, Aug 2014.

[11] Liang Tang, Yexi Jiang, Lei Li, and Tao Li, "Ensemble contextual bandits for personalized recommendation," in *Proceedings of the 8th ACM Conference on Recommender Systems*, New York, NY, USA, 2014, RecSys '14, pp. 73–80, ACM.

[12] Janis P. Hardwick, Quentin F. Stout, and Statistics Department, "Bandit strategies for ethical sequential allocation," *Comp. Sci. and Statist*, pp. 421–424, 1991.

[13] Y. Gai, B. Krishnamachari, and R. Jain, "Learning multiuser channel allocations in cognitive radio networks: A combinatorial multi-armed bandit formulation," in *New Frontiers in Dynamic Spectrum, 2010 IEEE Symposium on*, April 2010, pp. 1–9.

[14] L. Lai, H. Jiang, and H. V. Poor, "Medium access in cognitive radio networks: A competitive multi-armed bandit framework," in *2008 42nd Asilomar Conference on Signals, Systems and Computers*, Oct 2008, pp. 98–102.

[15] Tyler Lu, Dávid Pál, and Martin Pál, "Contextual multi-armed bandits.," in *AISTATS*, 2010, pp. 485–492.

[16] Tze Leung Lai, Philip W. Lavori, and Ka Wai Tsang, "Adaptive design of confirmatory trials: Advances and challenges," *Contemporary Clinical Trials*, vol. 45, Part A, pp. 93 – 102, 2015, 10th Anniversary Special Issue.

[17] K. Liu and Q. Zhao, "Distributed learning in multi-armed bandit with multiple players," *IEEE Transactions on Signal Processing*, vol. 58, no. 11, pp. 5667–5681, Nov 2010.

[18] Michel Tokic, "Adaptive $\epsilon$-greedy exploration in reinforcement learning based on value differences," in *Berlin / Heidelberg*. 2010, pp. 203–210, Springer.

[19] Olivier Chapelle and Lihong Li, "An empirical evaluation of thompson sampling," in *Advances in neural information processing systems*, 2011, pp. 2249–2257.

[20] John Langford and Tong Zhang, "The epoch-greedy algorithm for multi-armed bandits with side information," in *Advances in neural information processing systems*, 2008, pp. 817–824.

[21] Liang Tang, Romer Rosales, Ajit Singh, and Deepak Agarwal, "Automatic ad format selection via contextual bandits," in *Proceedings of the 22nd ACM international conference on Conference on information & knowledge management*, New York, NY, USA, 2013, CIKM '13, pp. 1587–1594, ACM.

[22] Peter Auer, Nicol Cesa-bianchi, Yoav Freund, and Robert E. Schapire, "The nonstochastic multiarmed bandit problem," *SIAM Journal on Computing*, vol. 32, pp. 2002, 2002.

[23] N. Cesa-Bianchi, Y. Freund, D. Haussler, D. P. Helmbold, R. E. Schapire, and M. K. Warmuth, "How to use expert advice," *Journal of the ACM*, vol. 44, no. 3, pp. 427–485, 1997.

[24] Alina Beygelzimer, John Langford, Lihong Li, Lev Reyzin, and Robert E Schapire, "Contextual bandit algorithms with supervised learning guarantees.," in *AISTATS*, 2011, pp. 19–26.

[25] Daniel Hsu, COLUMBIA EDU, and Robert E Schapire, "Taming the monster: A fast and simple algorithm for contextual bandits," 2014.

[26] Miroslav Dudik, Daniel Hsu, Satyen Kale, Nikos Karampatziakis, John Langford, Lev Reyzin, and Tong Zhang, "Efficient optimal learning for contextual bandits," in *Proceedings of the Twenty-Seventh Conference Annual Conference on Uncertainty in Artificial Intelligence (UAI-11)*, Corvallis, Oregon, 2011, pp. 169–178, AUAI Press.

[27] Peter Auer, "Using confidence bounds for exploitation-exploration tradeoffs," *Journal of Machine Learning Research*, vol. 3, no. Nov, pp. 397–422, 2002.

[28] Alekh Agarwal, Miroslav Dudík, Satyen Kale, John Langford, and Robert E Schapire, "Contextual bandit learning with predictable rewards.," in *AISTATS*, 2012, pp. 19–26.

[29] Frans MJ Willems, Yuri M Shtarkov, and Tjalling J Tjalkens, "The context-tree weighting method: basic properties," *IEEE Transactions on Information Theory*, vol. 41, no. 3, pp. 653–664, 1995.

[30] Kunihiko Sadakane, Takumi Okazaki, and Hiroshi Imai, "Implementing the context tree weighting method for text compression," in *Data Compression Conference, 2000. Proceedings. DCC 2000*. IEEE, 2000, pp. 123–132.

[31] Imre Csiszár and Zsolt Talata, "Context tree estimation for not necessarily finite memory processes, via bic and mdl," *IEEE Transactions on Information theory*, vol. 52, no. 3, pp. 1007–1016, 2006.

[32] T. Dumont, "Context tree estimation in variable length hidden markov models," *IEEE Transactions on Information Theory*, vol. 60, no. 6, pp. 3196–3208, June 2014.

[33] Fulvio Babich, Owen E Kelly, and Giancarlo Lombardi, "A context-tree based model for quantized fading," *IEEE communications letters*, vol. 3, no. 2, pp. 46–48, 1999.

[34] Suleyman S Kozat, Andrew C Singer, and Georg Christoph Zeitler, "Universal piecewise linear prediction via context trees," *IEEE Transactions on Signal Processing*, vol. 55, no. 7, pp. 3730–3745, 2007.

[35] N Denizcan Vanli and Suleyman Serdar Kozat, "A comprehensive approach to universal piecewise nonlinear regression based on trees," *IEEE Transactions on Signal Processing*, vol. 62, no. 20, pp. 5471–5486, 2014.

[36] H. Ozkan, N. D. Vanli, and S. S. Kozat, "Online classification via self-organizing space partitioning," *IEEE Transactions on Signal Processing*, vol. 64, no. 15, pp. 3895–3908, Aug 2016.

[37] Shai Shalev-Shwartz, "Online learning and online convex optimization," *Found. Trends Mach. Learn.*, vol. 4, no. 2, pp. 107–194, Feb. 2012.

[38] Lihong Li, Wei Chu, John Langford, and Xuanhui Wang, *Unbiased Offline Evaluation of Contextual-bandit-based News Article Recommendation Algorithms*, WSDM '11. ACM, New York, NY, USA, 2011.

[39] Ian Jolliffe, *Principal component analysis*, Wiley Online Library, 2002.